arXiv:1604.02388v2 [cs.CV] 9 Jun 2016

RGBD Semantic Segmentation Using Spatio-Temporal Data-Driven Pooling

Yang He[1], Wei-Chen Chiu[1], Margret Keuper[2] and Mario Fritz[1]

[1]Max Planck Institute for Informatics [2]University of Freiburg

Abstract. Beyond the success in classification, neural networks have recently shown strong results on pixel-wise prediction tasks like image semantic segmentation on RGBD data. However, the commonly used deconvolutional layers for upsampling intermediate representations to the full-resolution output still show different failure modes, like imprecise segmentation boundaries and label mistakes in particular on large, weakly textured objects (e.g. fridge, whiteboard, door). We attribute these errors in part to the rigid way, current network aggregate information, that can be either too local (missing context) or too global (inaccurate boundaries). Therefore we propose a data-driven pooling layer that integrates with fully convolutional architectures and utilizes boundary detection from RGBD image segmentation approaches. We extend our approach to leverage region-level correspondences across images with an additional temporal pooling stage. We evaluate our approach on the *NYU–Depth–V2* dataset comprised of indoor RGBD video sequences and compare it to various state-of-the-art baselines. Besides a general improvement over the state-of-the-art, our approach shows particularly good results in terms of accuracy of the predicted boundaries and in segmenting previously problematic classes.

1 Introduction

Consumer friendly and affordable combined image and depth-sensors such as *Kinect* are nowadays commercially deployed in scenarios such as gaming, personal 3D capture and robotic platforms. For such data, semantic image segmentation is an active research field. The aim is to provide from a set of known classes a correct label for each pixel in the image plane. Semantic segmentation has broad applicability and achieved great progress in the past year, in particular due to deep neural networks. Most recently, fully convolutional architectures [1] (FCN) have been proposed for this task, aiming at an end-to-end, joint learning from RGBD data to semantic segmentations. Albeit its success in providing quality estimates of semantic labels on coarse regions, FCN is at its core based on a rather rigid scheme with fixed rectangular pooling and upsampling operators. These lead to a fixed receptive field structure, which is negligent of the observed data. The consequences we have observed are twofold. On the one hand, the nature of the subsampling stages in the architecture of FCNs can cause a loss of spatial resolution in the deep layers. The final predictions, computed from upsampling low-resolution segmentations by the deconvolutional layers, apparently

lack the details necessary for accurate boundary localization due to too coarse inference schemes (large receptive field). On the other hand, we have observed a weakness of such methods on larger, less structured areas that belong to label classes such as "clothes", "curtain", "book", "whiteboard". We hypothesize that for such classes this inference scheme is too fine grained (small receptive field) in order to get enough context for successful labeling.

This motivates us to propose a data-driven pooling scheme, allowing to leverage highly accurate non-semantic boundary information in CNN based semantic segmentation. In addition to leveraging spatial information, the proposed scheme facilitates pooling over time or varying viewpoints. For many scenarios of practical relevance such as robotics, an image sequence is naturally collected and provides a substantially richer source of information than a single image. With additional evidence from the images on the same scene recorded from a different view, we can get more information of the objects in an image, such as different viewpoints or scales. These cues potentially improve the accuracy of semantic segmentations.

The main contributions of our paper are:

- Instead of refining the boundaries in the semantic segmentation results of FCN post hoc by using CRF models [2] or learning semantic region similarity [3], we build on advances in RGBD image segmentation approaches [4] which already show accurate boundary detections and provide a guide for the inference. A data-driven pooling scheme is proposed that uses superpixels generated by RGBD image segmentation for defining the receptive field of pooled pixel-label predictions.
- Our scheme naturally extends to spatio-temporal pooling. The correspondence between the superpixels across frames is built using optical flow. Then, a pooling operator is proposed on the connected segments to generate the final segment labeling.
- Our proposed method is evaluated on the challenging semantic segmentation dataset *NYU–Depth–V2* and outperforms several baselines as well as the state-of-the-art on the dataset. In particular, we improve on boundary precision and accuracy, as well as on difficult classes not well captured by other methods.

We organize the upcoming content of the paper as: In section 2, we first review the related work on convolutional networks based semantic segmentation, its extensions, and semantic segmentation based on multiple frames as observation. We depict our approach in section 3, and present the experimental results on *NYU–Depth–V2* [5] in section 4. Finally, this paper is concluded in section 5.

2 Related work

2.1 Fully convolutional networks and extensions:

Fully convolutional networks (FCN) [1], built on deep classification networks [6,7], carried their success forward to semantic segmentation networks that are

end-to-end trainable. FCN uses raw pixels as input and computes hierarchical visual features from a pretrained model. Finally, upsampling layers are used to directly generate dense predictions at the right resolution. Several extensions have been proposed to the standard FCN. Chen *et al.* combined the strengths of conditional random field (CRF) and FCN to refine the prediction, thus achieving more favorable results. However, the CRF and the network are trained separately in their work. Zheng *et al.* formulated CRFs as recurrent neural networks (RNN), and trained the FCN and their CRF-RNN end-to-end. Noh *et al.* [8] showed that FCN has difficulties to extract small objects because the method [1] only uses one deconvolution layer with large kernel size, which neglects a lot of useful information. Thus, they learned a network with multiple deconvolution layers to upsample the activations of the FCN to achieve more accurate and robust segmentations.

2.2 Boundary for Semantic Segmentation:

Boundary probabilities and edge detections have been exploited in many prior works on segmentation. Specifically, boundary information has recently been used in semantic segmentation by combining deep neural networks. Dai *et al.* [9] designed a convolutional feature masking layer for semantic segmentation, which allows networks to extract features in a stuff region with the help of superpixels. Gadde *et al.* [10] improved the semantic segmentation using superpixel convolutional networks with bilateral inception, which can control to merge initial superpixels by parameters and generate different levels of regions. Chen *et al.* [11] proposed a network to predict task-specific edges and semantic segmentations simultaneously. As our approach, this method is built on the FCN model, and refines the output of the FCN by the estimated edges. In the above mentioned methods [10,11], segmentation is improved by using estimated edges or pre-computed superpixels. Both methods showed the merit of providing boundary information to networks, which can generate more accurate segmentations. Our work follows a similar idea, but expresses the combination of boundaries and semantic segmentation as a data-driven pooling strategy which still allows end-to-end training and a natural extension to spatio-temporal pooling.

2.3 Multi-view semantic segmentation for indoor scenes:

The aim of multi-view semantic segmentation is to obtain better segmentations than methods relying on single view information. Couprie *et al.* [12] performed single image semantic segmentation with learned features with color and depth information, and applied a temporal smoothing in test time to improve the performance of frame-by-frame estimations. Hermans *et al.* [13] used the Bayesian update strategy to fuse new classification results and a CRF model in 3D space to smooth the segmentation. Stückler *et al.* [14] developed a real-time system, which use random forests to predict single view segmentations, and fuse all views to final output by simultaneous localization and mapping for establishing correspondence. In contrast to the above methods, we utilize optical flow and image superpixels to establish region level correspondences, and design an end-to-end trainable network which introduces correspondences between multiple frames.

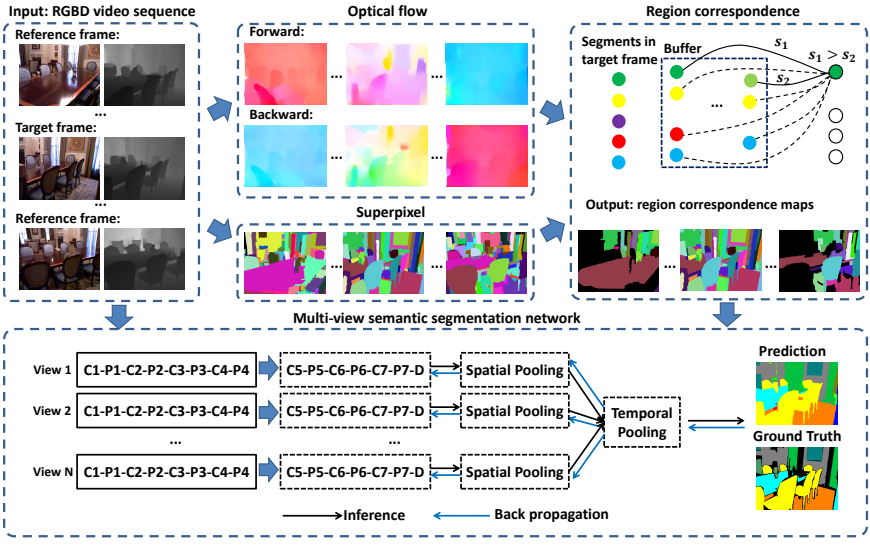

Fig. 1. Pipeline of our method. Our multi-view semantic segmentation network is built on the FCN ("C" denotes convolution layer, "P" denotes pooling layer, and "D" denotes the deconvolution layer). Our network fuses the information from multiple views with spatial-temporal data-driven pooling.

3 Approach

To overcome the limitations of a fixed receptive field, we propose to leverage segmentations (not class specific) in a data-driven pooling scheme. Hereby, we draw on the rich body of work on image segmentation and combine it with recent deep learning approaches to semantic image segmentations with accurate boundaries and labels. Our approach naturally extends to image sequences in which we can employ spatio-temporal pooling for improved semantic segmentation performance.

3.1 Overview

Our method takes an image sequence as input. We are interested in providing the most accurate semantic segmentation of one view in the sequence, which we call *target frame*. Other frames are called *reference frames*. Generally, the target frame is located at an arbitrary position in an image sequence. We leverage additional information from reference frames to improve the performance in the target frame. There are two main aspects in our system: the region level correspondence and a multi-view semantic segmentation network with data-driven pooling.

The pipeline of our method is summarized in Fig. 1. We first compute the superpixel segmentation of each frame and establish region level correspondences. Via the proposed data-driven pooling, information is aggregated in a deep learning architecture that is imposed by the correspondence structure. Hereby, we

achieve a tight integration of the superpixel segmentation into a deep learning framework that still allows for end-to-end training.

3.2 Region Level Correspondence

Mapping information from reference frames to the target frame is central to our framework. Here, we assume that, first, not all objects in the target frame are necessarily visible in every reference frame and secondly, for some areas it is easier to establish correspondences than for others. Therefore, video supervoxel methods such as [15] that force interframe correspondences and do not offer any confidence measure are not suitable. Instead, we establish the required correspondences on a frame-wise region level. In semantic segmentation, using regions has already proven to be effective in [10,16,17]. We use optical flow to establish the required correspondences between different views in a sequence. Ideally, these precise pixel level correspondences could directly support the segmentation of the target frame. Yet, it is not trivial to decide what should happen with points that have been occluded in previous frames or will undergo occlusion in subsequent frames. Additionally, errors in the optical flow can accumulate over time. However, entire regions are unlikely to undergo occlusion and single errors in the optical flow have little overall impact on a region level [18]. A further advantage is that it is more efficient to reject entire regions in reference frames if we cannot find a sufficiently confident correspondence to the target frame than it would be to make these decisions on a pixel level.

Superpixels To partition an RGBD image into several regions, we compute RGBD superpixels [4] in each frame, and optical flow between each pair of consecutive frames. Epic flow [19] is an edge-perserving optical flow technique, which uses boundary probabilities as an input. To take advantage of the depth information, we utilize the RGBD version of the structured edge detection [20] to generate boundary estimates. Then Epic flow is computed in forward and backward directions.

Robust Spatio-Temporal Matching The goal of a region correspondence is to project reliable information of reference frames to the target frame. To increase the probability of matching a region to the target frame, we set a buffer, as shown in Fig. 1, storing the superpixels of several previous frames to establish correspondences for the current frame. For pairs of regions between reference and target frames, we compute the intersection over union as a matching score. Assume R_c is a region in the current frame and R_p is a candidate region in the previous frame. With forward optical flow, we warp the region R_p to $R_p^{'}$, and compute the intersection over union between R_c and $R_p^{'}$, denoted as $\overrightarrow{IoU}_{cp}$. Similarly, we compute $\overleftarrow{IoU}_{cp}$ with backward optical flow. We regard R_c and R_p as a successful matching pair if their matching score meets $min(\overleftarrow{IoU}_{cp}, \overrightarrow{IoU}_{cp}) > threshold$. If R_c has several matching pairs, we only keep the one with the highest matching score. Keeping only one correspondence is effective as well as efficient. We will discuss in how far this is also sufficient in section 4.1. After matching all regions, we obtain a region correspondence map for each frame, which will be fed into the network.

3.3 Multi-view Semantic Segmentation Network with Data-Driven Pooling

The power of CNNs has been exploited in many computer vision tasks [6,21,22,23] by using a pooling technique to extract holistic or localized features for whole images[6], or proposals[22]. Most of them extract features in a rectangular region with pre-defined size, which fixes the aperture and therefore also the receptive field. To adapt the receptive field to the data and to aggregate the information from multiple views (the number of views varies for different regions), we need a flexible pooling strategy to train the whole model with multiple frames end-to-end. Thus we propose a data-driven pooling mechanism which extracts high-level features guided by superpixels.

Data-Driven Pooling We aim at adapting the pre-trained FCN-VGG16 to perform multi-view segmentation with region level correspondences. To this end, we include a data-driven pooling layer after the last deconvolution layer of the FCN, as shown in the bottom of Fig. 1. We only finetune the layers after *pool4* layer, as previous results have shown that layers down to *pool4* are quite stable [1]. Hence, we cache the output of the *pool4* layer, and train a networks taking *pool4* as input.

We refine the output of the deconvolution layer with superpixels and aggregate the information from multiple views by our proposed data-driven pooling, which can be decomposed into two steps: spatial pooling and temporal pooling. We can apply average pooling or max pooling to incorporate the superpixels. For max pooling, the maximum number within each region or temporal correspondence is mapped to the output side. In this section, we elaborately formulate the forward and backward propagation of the proposed data-driven pooling layer with average pooling.

The spatial pooling takes as input a feature map $I_s \in R^{N\times C\times H\times W}$ and a superpixel map $S \in R^{N\times H\times W}$, and generates the output $O_s \in R^{N\times C\times H\times W}$. Each channel of the frames is computed separately. The superpixel map S guides the forward and backward propagation of the layer. Here, $\Omega_{ij} \in \{(x,y)|S(i,x,y) = j\}$ denotes a superpixel in the i-th frame with segment index j. Then, the forward propagation of spatial average pooling can be formulated as

$$O_s(i,c,x_t,y_t) = \frac{1}{|\Omega_{ij}|}\sum_{(x_k,y_k)\in\Omega_{ij}} I_s(i,c,x_k,y_k) \tag{1}$$

for each channel index c and segment index j, where $(x_t,y_t) \in \Omega_{ij}$. Thus the posteriors of all pixels in the j-th region of the i-th frame are enforced to be the same. To train our model, we employ stochastic gradient descent for optimization. The gradient of the input $I(i,c,x_k,y_k)$, where $(x_k,y_k) \in \Omega_{ij}$, in spatial average pooling is calculated by back propagation [24],

$$\begin{aligned}\frac{\partial L}{\partial I_s(i,c,x_k,y_k)} &= \sum_{(x_t,y_t)\in\Omega_{ij}} \frac{\partial L}{\partial O_s(i,c,x_t,y_t)} \frac{\partial O_s(i,c,x_t,y_t)}{\partial I_s(i,c,x_k,y_k)} \\ &= \frac{1}{|\Omega_{ij}|} \sum_{(x_t,y_t)\in\Omega_{ij}} \frac{\partial L}{\partial O_s(i,c,x_t,y_t)}.\end{aligned} \tag{2}$$

Similarly, we formulate the temporal average pooling which fuses the information from multiple frames $I_t \in R^{N\times C\times H\times W}$ to one frame $O_t \in R^{C\times H\times W}$. Besides the superpixel map $S \in R^{N\times H\times W}$, this layer also needs a superpixel map on the target frame $S_{target} \in R^{H\times W}$, where $\widetilde{\Omega}_j \in \{(x,y)|S_{target}(x,y)=j\}$ is the superpixel with index j in the target frame. The forward propagation can be expressed as:

$$O_t(c,x_t,y_t) = \frac{1}{K}\sum_{i=1}^{N} \sum_{(x_k,y_k)\in\Omega_{ij}} \frac{1}{|\Omega_{ij}|} I_t(i,c,x_k,y_k), \tag{3}$$

where $(x_t,y_t)\in\widetilde{\Omega}_j$, Ω_{ij} is the superpixel with index j of the i-th input frame, and $K = |\{|\Omega_{ij}|>0|1\le i\le N\}|$. The term $1/|\Omega_{ij}|$ is used to eliminate the effect of differences in region size. The gradient is calculated by

$$\begin{aligned}\frac{\partial L}{\partial I_t(i,c,x_k,y_k)} &= \sum_{(x_t,y_t)\in\widetilde{\Omega}_j} \frac{\partial L}{\partial O_t(c,x_t,y_t)} \frac{\partial O_t(c,x_t,y_t)}{\partial I_t(i,c,x_k,y_k)} \\ &= \frac{1}{K|\Omega_{ij}|} \sum_{(x_t,y_t)\in\widetilde{\Omega}_j} \frac{\partial L}{\partial O_t(c,x_t,y_t)}.\end{aligned} \tag{4}$$

Implementation Details We regard the frames with groundtruth annotations as target frames. For each target frame, we equidistantly sample up to 100 frames around it with the static interval of 3 frames. In results, we obtain 1449 RGBD sequences of 28∼101 frames. Next, we compute the superpixels [4] and Epic flow [19] with the default settings provided in the corresponding source codes. The buffer size and the threshold are 3 and 0.4 for the computation of region correspondences (cf. section 3.2). Finally, for each RGBD sequence, we randomly sample 11 frames including the target frame together with their correspondence maps as the input for our network. We implement the proposed network using the *Caffe* framework [25]. We use RGB images and HHA representations of depth and train the network by stochastic gradient descent with momentum term. Due to the memory limitation, we first run FCN-VGG16 and cache the output *pool4_rgb* and *pool4_hha* and then finetune the layers after *pool4* with a new network which is the copy of higher layers in FCN. We use a minibatch size of 10, momentum 0.9, weight decay 0.0005 and fixed learning rate 10^{-14}. We finetune our model by using cross entropy loss with 1000 iterations for all our models in the experiments, which decrease the training loss approximately from 2×10^5 to less than 10^5.

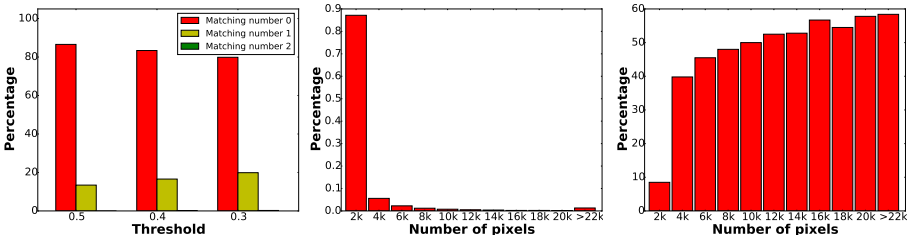

Fig. 2. Statistics on region correspondence. The left figure draws the histogram of successful matching number of regions between adjacent frames with different thresholds. The middle figure draws the histogram of region size. The right figure draws the histogram of matching number of regions in the whole video sequence.

4 Experiments and Analysis

The experiments are performed on the *NYU–Depth–V2* (NYUDv2) dataset, with various settings of groundtruth annotations: 4 classes [5], 13 classes [12], and 40 classes [26]. The NYUDv2 dataset contains 518 RGBD videos, which have more than 400,000 images. Among them, there are 1449 densely labeled frames, which are split into 795 training images and 654 testing images. We report the results on the labeled frames, using the same evaluation protocol and metrics as [1], pixel accuracy (*Pixel Acc.*), mean accuracy (*Mean Acc.*), region intersection over union (*Mean IoU*), and frequency weighted intersection over union (*f.w. IoU*).

4.1 Region Correspondence Analysis

As establishing good correspondence between reference and target frames is important in our method, we start with a study of the computed regions and their correspondence. We first make statistics on how many regions can be matched for a region in its neighboring frames as shown in Fig. 2 (left). We observe that most regions cannot find any match in other frame. In order to understand this phenomenon, we count the size of regions provided by superpixelization in all 1449 labeled frames, as shown in Fig. 2 (middle). Besides, we count the number of overall matching regions for the segments in target frames, and get the distribution by the size of regions illustrated in Fig. 2 (right). It shows that most regions are small (less than 2000 pixels), and they are unlikely to find good matching regions. That is why there are almost 80% of the segments without correspondence. Consequentially, most regions with matches are with relatively bigger size and cover large portion of images, and they usually have only one match to their neighboring frames. Therefore, we decide to only consider the correspondence with the highest score for our data-driven pooling method which also significantly reduces the computational cost.

4.2 Average vs. Max Data-Driven Spatio-Temporal Pooling

Our data-driven pooling aggregates the local information from multiple pixels within a segment and across multiple views. Average pooling and max pooling

Table 1. Comparison results of average and max data-driven spatio-temporal pooling.

Segment	Multi-View	Pixel Acc.	Mean Acc.	Mean IoU	f. w. IoU
Average	Average	70.1	53.8	40.1	55.7
Average	Max	69.4	51.0	38.0	54.4
Max	Average	66.4	45.4	33.8	49.6
Max	Max	64.9	44.5	32.1	47.9

Table 2. The performance of oracle case using groundtruth to label the regions.

Groundtruth	Pixel Acc.	Mean Acc.	Mean IoU	f. w. IoU
Current Frame	96.2	94.0	90.2	92.7
Next Frame	84.7	76.2	63.4	74.4

are canonical choices used in many deep neural network architectures that allow for end-to-end training. Here we test the average pooling and max pooling both in spatial and temporal pooling and show the results in Table 1. All the models are trained with multiple frames, and tested on multiple frames. Specifically, the first column specifies the spatial pooling type over the segment and the second column specifies the temporal pooling type across frames. Average pooling turns out to perform best for spatial and temporal pooling and therefore we use this combination in the rest of our experiments.

4.3 Oracle Performance using Groundtruth Labels

We perform two best-case analysis by computing an oracle performance where groundtruth labels are available for either reference or target frames. The first row of Table 2 shows the achievable performance by performing a majority vote of the groundtruth pixel labels on the employed superpixels from [4]. Thereby we achieve an upper bound of 96.2% on the pixel accuracy that is implied by the superpixel over-segmentation. In order to evaluate the effectiveness of our region correspondence, we use groundtruth labels of reference frames in the sequence. We collect 143 views to conduct this experiment in NYUDv2, which have corresponding regions in target frames. We ignore regions without correspondence in the next frame to compute the quantitative results, which are presented in Table 2. This best-case analysis for correspondence results in a pixel accuracy of 84.7%. Both oracle performances indicate a strong potential for performance improvements in our setup in all 4 reported measures. In addition, the bottom two rows in Fig. 3 show examples labeling for the best-case study.

4.4 Groundtruth Analysis

At a closer look, it turns out that at least part of the performance loss in the best-case analysis for the correspondence is not due to bad matches between regions. In Fig. 3, we present some examples of the annotations provided in the dataset. In several cases, as the ones shown in the figure, the labeling is inconsistent and object labels are changed during the sequence. From left to right in Fig. 3, table changes to desk, table changes to dresser, floor changes to floor mat, bookshelf changes to shelves, cabinet changes to other-furniture, and window changes to

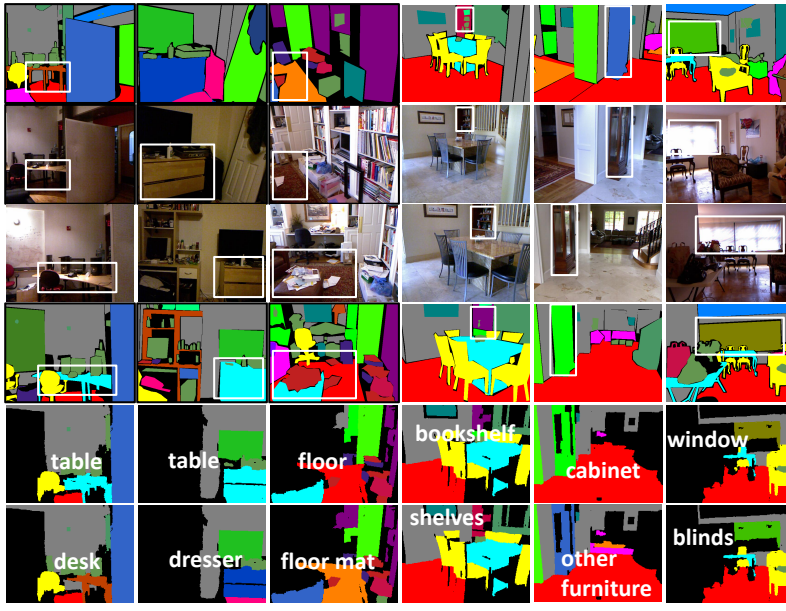

Fig. 3. Example of groundtruth limitation and segmentation results of oracle case. Row 3 and 2 draw color images of target frame and next labeled frame, respectively. And row 4 and 1 draw their groundtruth. The segmentation result with groundtruth of target frame is shown in row 5, and the result with groundtruth of next frame is shown in row 6. We point out the regions in different frames with white bounding box, which are the same object of different views but labeled as different classes.

blinds. Consequently, we see mistakes in the last two rows corresponding to the best case results due to inconsistent labelings.

4.5 Analysis of Multi-View Prediction

In our multi-view model, we subsample frames from the whole video for computational considerations. There is a trade-off between close-by and distant frames to be made. If we select frames far away from the target frames, they can provide more diverse views of an object, while matching is more challenging and potentially less accurate than for close-by frames. Hence, we analyze the influence of the distance of selected frames to target frames, and report the *Mean Acc.* and *Mean IoU* in Fig 4. In results, providing wider views is helpful, as the performance is improved with the increase of max distance. And selecting the data in the future, which is another way to provide wider views, also contributes to the improvements of performance.

4.6 Comparison to Baseline Methods

We compare to three baselines: fully convolutional networks (FCN) [1], multi-view pixel-level correspondence network, and single-view spatial pooling net-

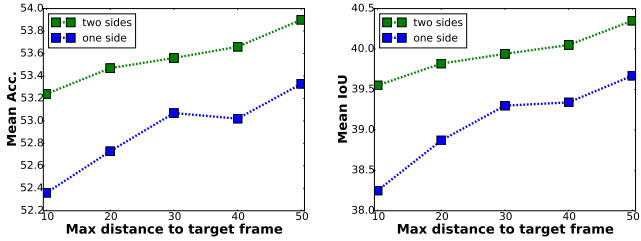

Fig. 4. The performance of multi-view prediction with different max distance. Green lines show the results of using the views in the future and past. Blue lines show the results of only using the past views.

work. The second baseline and the third baseline are built on the top of FCN-VGG16. These baselines aim at analyzing the performance of the three key ingredients that we are building on: FCN, flow, superpixels. The quantitative comparison is listed in Table 3.

The second baseline, multi-view pixel-level correspondence network (Multi-view Pixel Network) uses the per-pixel correspondence by Epic flow, and applies average pooling to fuse the information from multiple view. Table 3 shows that baseline 2 has already better performance than FCN in 3 out of 4 measures, which means leveraging multiple views benefits the estimation of a single pixel. But the optical flow is not perfect, thus obtaining accurate pixel-level correspondence is challenging. Consequently, the full model, which is built on the region-level correspondence, has a significant improvement over this baseline and FCN in all 4 measures.

The third baseline introduces superpixel to the network (Single-view Superpixel Network), and we train and test it using only one view. The data-driven pooling scheme divides the FCN prediction into several segments and refines the prediction by enforcing the pixels in a segment having the same posterior.

Table 3. Comparison results with baselines on NYUDv2

Methods	Pixel Acc.	Mean Acc.	Mean IoU	f. w. IoU
FCN [1]	65.4	46.1	34.0	49.5
Multi-view Pixel Network	66.2	45.9	34.6	50.2
Single-view Superpixel Network	68.5	48.7	36.0	52.9
Our full model	**70.1**	**53.8**	**40.1**	**55.7**

Table 4. Performance on the 4-class (left) and 13-class (right) semantic segmentation tasks.

	Pixel Acc.	Mean Acc.	Pixel Acc.	Mean Acc.
Couprie *et al.* [12]	64.5	63.5	52.4	36.2
Hermans *et al.* [13]	69.0	68.1	54.2	48.0
Stückler *et al.* [14]	70.6	66.8	-	-
Eigen *et al.* [27]	<u>83.2</u>	<u>82.0</u>	<u>75.4</u>	66.9
Ours (Single-View)	82.7	81.3	74.8	<u>67.0</u>
Ours (Multi-View)	**83.6**	**82.5**	**75.8**	**68.4**

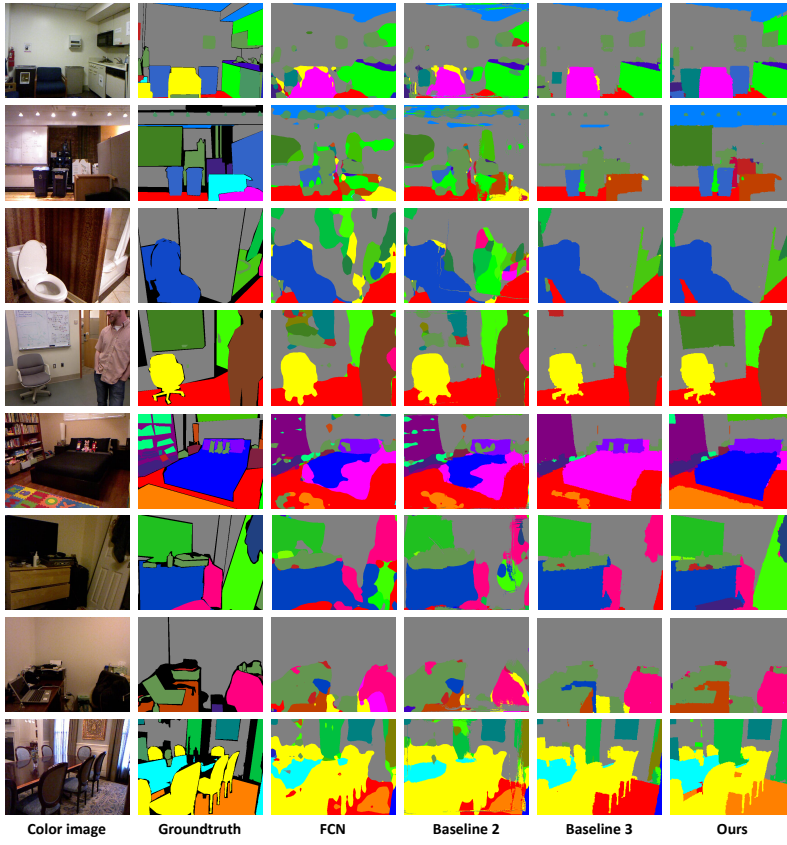

Fig. 5. Visualization examples of semantic segmentation on NYUDv2. Column 1 shows the RGB images and column 2 shows the ground truth (black represents the unlabeled pixels).column 3 shows the results from FCN [1], column 4 and 5 shows the results from the second baseline and the third baseline, and the results from our whole system are shown in the column 6. Best view in color.

For each segment, we apply average pooling to extract the representation of the segments. Comparing baseline 3 with FCN, we observe a performance improvement on all four metrics of at least 2 percentage points (*pp*), which clearly shows the effectiveness of combining superpixel with convolutional networks. Our full model consistently improves over all baselines and outperforms the FCN model $4.7pp, 7.7pp, 6.1pp, 6.2pp$.

Besides, we also report the qualitative results as shown in Fig 5. Our data-driven pooling networks can generate more precise and smooth segmentation than FCN and multi-view networks with pixel correspondence. And it also achieves more accurate estimation than single-view model. Therefore, leveraging multi-view data and using data-driven pooling to integrate the spatial-temporal information are effective.

4.7 Analysis of Semantic Segmentation Boundary Accuracy

In order to quantify the improvement on semantic boundary localization based on the proposed data-driven pooling scheme, we use Boundary Precision Recall (BPR), as also used in image or video segmentation benchmark [28,29] for evaluation. Fig 6 shows the resulting semantic boundary average precision-recall curve. We conclude that our method generates more accurate boundaries than FCN, which achieve 0.477 BPR score while our method achieves 0.647. Besides, our method even improves on the superpixel [4] we build on, which means our method can successfully merge over-segmentations or non-semantic boundaries between adjacent instances of the same semantic class.

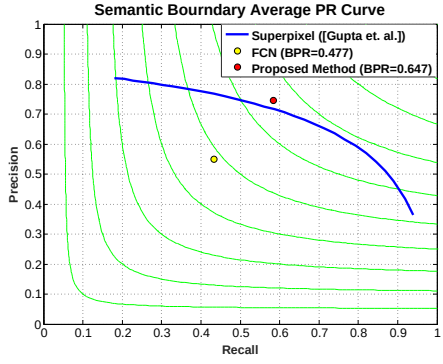

Fig. 6. Precision-recall curve on semantic boundaries on the NYUDv2 dataset.

4.8 Comparison with State-of-the-Art

We first evaluate and analyze our methods on the NYUDv2 40-class task and compare our method to state-of-the-art methods [1,4,30,27,17]. The quantitative results are shown in Table 5. The multi-view model (Multi-View) leverages multiple views from unlabeled data in NYUDv2, and we also test our single-view model (Single-View) that only uses target frames. Already our single view results – which only use additional data at training time – are better than the state of the art performance in all 4 metrics by $2.9pp, 6.0pp, 4.0pp, 2.6pp$ respectively. Our multi-view model achieves a consistent improvement over the single-view model and outperform all competitors in all four metrics by $2.1pp, 7.7pp, 6.0pp, 4.3pp$ respectively. Particular strong improvements are observed on challenging objects such as whiteboard, bag, other furniture where we roughly double the performance. In particular, the substantial improvement on the class averaged accuracy is encouraging as here no particular fitting to single classes or large classes is observed.

Additionally, we report the performance numbers for 4-class and 13-class tasks, and compare our approach to previous state-of-the-art methods [12,13,14,27] in Table 4, including the multi-view methods [13,14]. We observe that our single-view model already outperforms the multi-view baselines. Our proposed multi-view pooling scheme further boosts *Pixel Acc.* and *Mean Acc.* by more than $1pp$ and outperforms the state-of-the-art [27]. In summary, our multi-view model achieves $0.4pp, 0.4pp, 4.5pp$ of *Pixel Acc.* and $0.5pp, 1.5pp, 8.7pp$ of *Mean Acc.* improvement with respect to [27] on 4-, 13- and 40-class tasks, which implies

Table 5. Performance on the 40-class semantic segmentation task. We compare our method with five state-of-the-art methods: [1,4,30,27] are also based on convolutional networks, and [17] is region labeling method, which is also related to ours. We mark the best performance in all methods with **BOLD** font, and the second best one is written with UNDERLINE.

	wall	floor	cabinet	bed	chair	sofa	table	door	window	bookshelf	picture	counter	blinds	desk	shelves
Long *et al.* [1]	69.9	79.4	50.3	66.0	47.5	53.2	32.8	22.1	39.0	36.1	50.5	54.2	45.8	11.9	8.6
Gupta *et al.* [4]	68.0	81.3	44.9	65.0	47.9	47.9	29.9	20.3	32.6	18.1	40.3	51.3	42.0	11.3	3.5
Kendall *et al.* [30]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Eigen *et al.* [27]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Deng *et al.* [17]	65.6	79.2	51.9	66.7	41.0	55.7	36.5	20.3	33.2	32.6	44.6	53.6	49.1	10.8	9.1
Ours (Single-View)	72.4	84.3	52.0	71.5	54.3	58.8	37.9	28.2	41.9	38.5	52.3	58.2	**49.7**	**14.3**	8.1
Ours (Multi-View)	**72.7**	**85.7**	**55.4**	**73.6**	**58.5**	**60.1**	**42.7**	**30.2**	**42.1**	**41.9**	**52.9**	**59.7**	46.7	13.5	**9.4**

	curtain	dresser	pillow	mirror	floormat	clothes	ceiling	books	fridge	tv	paper	towel	showercurtain	box	whiteboard
Long *et al.* [1]	32.5	31.0	37.5	22.4	13.6	18.3	59.1	27.3	27.0	41.9	15.9	26.1	14.1	6.5	12.9
Gupta *et al.* [4]	29.1	34.8	34.4	16.4	28.0	4.7	**60.5**	6.4	14.5	31.0	14.3	16.3	4.2	2.1	14.2
Kendall *et al.* [30]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Eigen *et al.* [27]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Deng *et al.* [17]	**47.6**	27.6	**42.5**	30.2	32.7	12.6	56.7	8.9	21.6	19.2	**28.0**	28.6	**22.9**	1.6	1.0
Ours (Single-View)	42.9	35.9	40.8	27.7	31.9	19.3	55.6	28.2	38.3	46.9	17.6	31.2	11.0	6.5	28.2
Ours (Multi-View)	40.7	**44.1**	42.0	**34.5**	**35.6**	**22.2**	55.9	**29.8**	**41.7**	**52.5**	21.1	**34.4**	15.5	**7.8**	**29.2**

	person	nightstand	toilet	sink	lamp	bathtub	bag	other struct	other furni	other props	Pixel Acc.	Mean Acc.	Mean IoU	f. w. IoU
Long *et al.* [1]	57.6	30.1	61.3	44.8	32.1	**39.2**	4.8	15.2	7.7	30.0	65.4	46.1	34.0	49.5
Gupta *et al.* [4]	0.2	27.2	55.1	37.5	34.8	38.2	0.2	7.1	6.1	23.1	60.3	-	28.6	47.0
Kendall *et al.* [30]	-	-	-	-	-	-	-	-	-	-	68.0	45.8	32.4	-
Eigen *et al.* [27]	-	-	-	-	-	-	-	-	-	-	65.6	45.1	34.1	51.4
Deng *et al.* [17]	9.6	30.6	48.4	41.8	28.1	27.6	0	9.8	7.6	24.5	63.8	-	31.5	48.5
Ours (Single-View)	**66.7**	34.1	**62.8**	**47.8**	35.1	26.4	**8.8**	**19.3**	10.9	29.2	68.4	52.1	38.1	54.0
Ours (Multi-View)	60.7	**42.2**	62.7	47.4	**38.6**	28.5	7.3	18.8	**15.1**	**31.4**	**70.1**	**53.8**	**40.1**	**55.7**

that recognizing an object or a region is easier from multiple views than single view, especially in fine-grained catergory settings.

5 Conclusion

In this paper, we have presented a novel semantic segmentation approach using image sequences. We design a multi-view semantic segmentation network with data-driven spatio-temporal pooling which can receive multiple images and their correspondence as input. We propagate the information from multiple views to the target frame, and significantly improve the semantic segmentation performance on the target frame. Besides, our method can leverage large scale unlabeled images for training, and we show that using additional data also benefits single image semantic segmentation.

Acknowledgments

We would like to thank Marcus Rohrbach for his support, and Alexander Hermans for providing groundtruth labels for the 13-class task on NYUd2. Margret Keuper was supported by an ERC Starting Grant VideoLearn.

References

1. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. (2015)
2. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected crfs. In: ICLR. (2014)
3. Harley, A.W., Derpanis, K.G., Kokkinos, I.: Learning dense convolutional embeddings for semantic segmentation. In: ICLR Workshop. (2015)
4. Gupta, S., Girshick, R., Arbeláez, P., Malik, J.: Learning rich features from rgb-d images for object detection and segmentation. In: ECCV. (2014)
5. Nathan Silberman, Derek Hoiem, P.K., Fergus, R.: Indoor segmentation and support inference from rgbd images. In: ECCV. (2012)
6. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. (2012)
7. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
8. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: ICCV. (2015)
9. Dai, J., He, K., Sun, J.: Convolutional feature masking for joint object and stuff segmentation. In: CVPR. (2015)
10. Gadde, R., Jampani, V., Kiefel, M., V. Gehler, P.: Superpixel convolutional networks using bilateral inceptions. arxiv (2015)
11. Chen, L.C., Barron, J.T., Papandreou, G., Murphy, K., Yuille, A.L.: Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. arxiv (2015)
12. Couprie, C., Farabet, C., Najman, L., LeCun, Y.: Indoor semantic segmentation using depth information. In: ICLR. (2013)
13. Hermans, A., Floros, G., Leibe, B.: Dense 3d semantic mapping of indoor scenes from rgb-d images. In: ICRA. (2014)
14. Stückler, J., Waldvogel, B., Schulz, H., Behnke, S.: Dense real-time mapping of object-class semantics from rgb-d video. Journal of Real-Time Image Processing **10**(4) (2015) 599–609
15. Grundmann, M., Kwatra, V., Han, M., Essa, I.: Efficient hierarchical graph based video segmentation. In: IEEE CVPR. (2010)
16. Arbeláez, P., Hariharan, B., Gu, C., Gupta, S., Bourdev, L., Malik, J.: Semantic segmentation using regions and parts. In: CVPR. (2012)
17. Deng, Z., Todorovic, S., Jan Latecki, L.: Semantic segmentation of rgbd images with mutex constraints. In: CVPR. (2015)
18. Brox, T., Malik, J.: Large displacement optical flow: descriptor matching in variational motion estimation. TPAMI (2011)
19. Revaud, J., Weinzaepfel, P., Harchaoui, Z., Schmid, C.: Epicflow: Edge-preserving interpolation of correspondences for optical flow. In: CVPR. (2015)

20. Dollár, P., Zitnick, C.: Structured forests for fast edge detection. In: CVPR. (2013)
21. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: Integrated recognition, localization and detection using convolutional networks. arXiv preprint arXiv:1312.6229 (2013)
22. Girshick, R.: Fast r-cnn. In: CVPR. (2015)
23. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. TPAMI (2015)
24. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. Cognitive modeling **5**(3) (1988) 1
25. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. In: Proceedings of the ACM International Conference on Multimedia. (2014)
26. Gupta, S., Arbelaez, P., Malik, J.: Perceptual organization and recognition of indoor scenes from rgb-d images. In: CVPR. (2013)
27. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. arxiv (2015)
28. Galasso, F., Nagaraja, N.S., Cardenas, T.J., Brox, T., Schiele, B.: A unified video segmentation benchmark: Annotation, metrics and analysis. In: ICCV. (2013)
29. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. TPAMI (2011)
30. Kendall, A., Vijay, B., Cipolla, R.: Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. arxiv (2015)